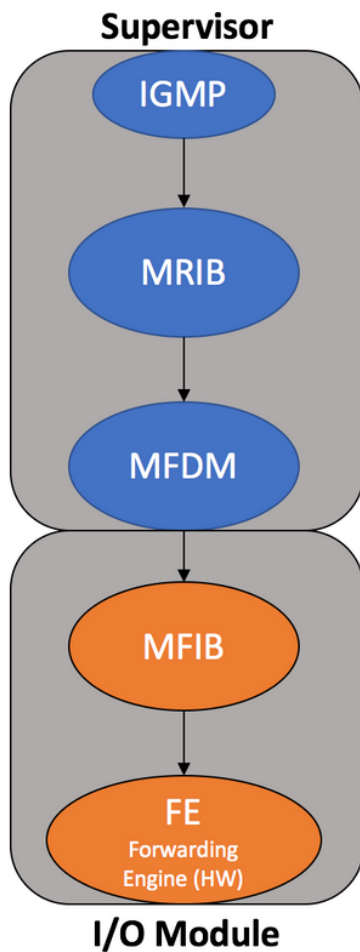


EXHIBIT 27

M3 Multicast Forwarding

This document covers both L2 and L3 multicast forwarding for the M3 module. It will include PI commands as well as PD commands specific to this module to verify hardware state.

L2 Multicast



PI

show ip mroute

M2RIB contains S,G for desired group and OIF as L2 VLAN that S resides in.

show ip igmp snooping vlan <x>

Check IGMP is enabled on given VLAN for interested receivers, ensure Active Ports contains expected L2 receivers.

Ensure a querier is present as is required for L2 multicast. Querier will be populated when N7K has "ip pim sparse-mode" enabled on SVI for the L2 VLAN. If no SVI, querier must be manually configured in VLAN configuration mode.

show ip igmp snooping groups vlan <x>

Check that expected L2 receiver ports are present next to VLAN, S, G.

show ip igmp route <group> <source> vlan <x>

Verify there is a last known reporter in the expected vlan/subnet. Check uptime is as expected.

show forwarding distribution ip igmp snooping vlan <x> group <y> detail

IGMP state in MFDM, ensure it contains expected L2 receiver ports in OIFL. Take note of platform index for future verification in HW.

show forwarding distribution l2 multicast vlan <x>

L2 mroutes in given VLAN present in MFDM on supervisor. Ensure OIFL is populated with expected L2 VLAN, and packet/byte counter periodically increments over time aligning to traffic Source is sending. Check platform index matches with IGMP MFDM state.

show forwarding distribution multicast route vrf <x> group <y>

L3 perspective, but useful for checking the packet/byte counter is increasing for the given S,G.

show forwarding vrf default multicast route group <x> source <y> module <z>

Verifying route is present and correct in MFIB on LC. Ensure OIFL is correct.(M3 does not have byte counter value, only packet).

show system internal forwarding l2 multicast vlan <x> module <y>

L2 mroutes in a given vlan present in L2MCAST software table on the given module. Check DTL value and swindex.

show system internal forwarding vrf default multicast route detail group <x> source <y>

HW programming state for the MFIB. Verify index's are valid and dest idx is not sending to a drop index.

show system internal pixm info ltl 0x

Use platform index from MFDM above. Ensure output contains expected L2 receiver ports. Note the MI value for next command. V5/V4 FPOE outputs should be noted for ELAM verification as well.

show system internal xbar static-mc

Check what groups expected egress module toward L2 receivers is present on. Use map to ensure group/module is present for the MI value found in PIXM for the L2 snooping index above.

PD**show system internal forwarding l2 multicast swindex swindex-tbl**

Per-instance (FE/SoC) sw_index to hw_index mapping for L2 multicast routes. This is HW state, verify BD and instance is correct, and sw-index matches from L2MCAST table output for the module. Note the HWPTR and MC_DI values.


show hardware internal forwarding l2 inst <x> table mac index <y>

L2 multicast routes programmed in MAC table hardware for given HWPTR index above.

show hardware internal forwarding l2 inst <x> table sw-mc-di index <y>

Example

Mod	Ports	Module-Type	Model	Status
1	0	Supervisor Module-2	N7K-SUP2E	active *
2	0	Supervisor Module-2	N7K-SUP2E	ha-standby
3	12	10/40 Gbps Ethernet Module	N7K-F312FQ-25	ok
4	24	10/40 Gbps Ethernet Module	N7K-M324FQ-25L	ok



IGMP

show ip igmp snooping vlan 1001

```
Global IGMP Snooping Information:
  IGMP Snooping enabled
  Optimised Multicast Flood (OMF) enabled
  IGMPv1/v2 Report Suppression enabled
  IGMPv3 Report Suppression disabled
  Link Local Groups Suppression enabled

IGMP Snooping information for vlan 1001
  IGMP snooping enabled
  Lookup mode: IP
  Optimised Multicast Flood (OMF) enabled
  IGMP querier present, address: 10.201.1.2, version: 2, i/f Vlan1001
  Querier interval: 125 secs
  Querier last member query interval: 1 secs
  Querier robustness: 2
  Switch-querier disabled
  IGMPv3 Explicit tracking enabled
  IGMPv2 Fast leave disabled
  IGMPv1/v2 Report suppression enabled
  IGMPv3 Report suppression disabled
  Link Local Groups suppression enabled
  Router port detection using PIM Hellos, IGMP Queries
  Number of router-ports: 2
  Number of groups: 2
  VLAN vPC function enabled
```

Active ports:
 Po2 Eth4/1/1 **Eth3/1/1** Eth102/1/9
 Eth101/1/9

show ip igmp snooping groups vlan 1001

Type: S - Static, D - Dynamic, R - Router port, F - Fabricpath core port

Vlan	Group	Address	Ver	Type	Port list
1001	*/*		-	R	Vlan1001 Po2
1001	233.1.0.1		v2	D	Eth3/1/1

show ip igmp route 233.1.0.1

IGMP Connected Group Membership for VRF "default" - 2 total entries

Type: S - Static, D - Dynamic, L - Local, T - SSM Translated

Group	Address	Type	Interface	Uptime	Expires	Last Reporter
233.1.0.1		D	Vlan1001	1d14h	00:02:58	10.201.1.23

MRIB

show ip mroute 233.1.0.1

IP Multicast Routing Table for VRF "default"

(**10.201.1.24/32**, **233.1.0.1/32**), uptime: 03:44:39, ip mrib pim
 Incoming interface: Vlan1001, RPF nbr: 10.201.1.24
 Outgoing interface list: (count: 2)
 Ethernet4/4, uptime: 03:44:39, pim
Vlan1001, uptime: 03:44:39, mrib, (RPF)

MFDM

show forwarding distribution ip igmp snooping vlan 1001 group 233.1.0.1 detail

Vlan: 1001, Group: 233.1.0.1, Source: 0.0.0.0
 Route Flags: 0
 Outgoing Interface List Index: 3
 Reference Count: 2
 Platform Index: **0x7fe4**
 Vpc peer link exclude flag set
 Number of Outgoing Interfaces: 2
 port-channel2
Ethernet3/1/1

show forwarding distribution l2 multicast vlan 1001

Vlan: 1001, Group: 233.1.0.1, Source: 0.0.0.0
 Route Flags: 0

```

Outgoing Interface List Index: 3
Reference Count: 2
Platform Index: 0x7fe4
Vpc peer link exclude flag set
Number of Outgoing Interfaces: 2
    port-channel2
    Ethernet3/1/1

```

show forwarding distribution multicast route vrf default group 233.1.0.1

```

(10.201.1.24/32, 233.1.0.1/32), RPF Interface: Vlan1001, flags:
Received Packets: 12662123 Bytes: 12961957636

```

PIXM

show system internal pixm info ltl 0x7fe4

```

MCAST LTLs allocated for VDC:3
=====
LTL      IFIDX/RID      LTL_FLAG CB_FLAG
0x7fe4  0x00000003  0x00      0x0002

mi | v5_f3_fpoe | v4_fpoe | v5_fpoe | clp_v4_l2 | clp_v5_l2 | clp20_v4_l3 |
clp_cr_v4_l3 | flag | proxy_if_index
0xc | 0xc | 0x5 | 0xc | 0x0 | 0xc | 0xc | 0xc | 0xc | 0x0 | none

Member info
-----
IFIDX      LTL
-----
Eth3/1/1      0x0000

```

show system internal xbar static-mc

```

Line card Module 3 groups: 1, 2, 3, 4
Line card Module 4 groups: 5, 6, 7, 8
-----
| Multicast Index | group-mask | List of groups(1-based) |
-----
<snip>
|           0012 | 0x000011 | 1,5 | <-- 0xc == 12

```

MFIB

show forwarding vrf default multicast route group 233.1.0.1 source 10.201.1.24 module 4

```
(10.201.1.24/32, 233.1.0.1/32), RPF Interface: Vlan100, flags:
Received Packets: 12662123 Bytes: 12961957636 <--- Packets hitting route
Number of Outgoing Interfaces: 2
Outgoing Interface List Index: 3
Vlan100 Outgoing Packets: 12662123 Bytes:0 <--- M3 has no byte counter for OIFs
```

FE (Hardware)

Supervisor - PI Level:

show system internal forwarding l2 multicast vlan 1001 module 4

Flag Type: R-Remote Receiver, L-Local Receiver, C-Copy-to-Sup Enabled, U-Undefined

Lookup Mode : IP

Vlan/SW_BD	BD	Ftag	Group	Source	RID	DTL	swindex	Fl
<snip>								
1001	43	0	233.1.0.1		3	0x7fe4	251	L

show system internal forwarding vrf default multicast route detail group 233.1.0.1 source 10.201.1.24

```
<snip>
slot 4
=====
```

Hardware Multicast FIB Entries:

Flags Legend:

```
* - s_star_priority
S - sg_entry
D - Non-RPF Drop
B - Bi-dir route W - Wildcard route
```

```
(10.201.1.24/32, 233.1.0.1/32), Flags: *S
```

```
Dev: 0, HWIndex: 0xfd6be DRAM Index:: 0xfd6be, VPN: 0x5 <--- Dev is FE instance,
VPN is VRF
```

```
RPF Interface: Vlan1001, LIF: 0x3e9
```

```
ML3 Adj/Rit Idx: 0x98/0x12f, INGRESS_MET: 0xe, EGRESS_MET: 0xb
```

```
PD oiflist Idx: 0x8
```

```
MD Adj/Rit Idx: 0x9c/0x8012, MDT Idx: 0x3, MTU Idx: 0x1, Dest Idx: 0x2865
```


```
Dev: 0 Adj/Rit Index: 0x290/0x8000
```

```
Type: OIF elif: 0x100a Ethernet4/4 <--- For L3 Mcast
```

```
Dest Idx: 0x42 SMAC: 6c9c.ed4b.aac3
```

Module - PD Level:

L3 Multicast



MRIB - (Multicast Routing Information Base) is responsible to cache routes added by the multicast- protocols - PIM, IGMP, MSDP, IP (each of which runs as separate process) and sends route- updates to MFDM (Multicast Forwarding Distribution Manager) to program the hardware multicast forwarding tables.

show ip mroute

```
IP Multicast Routing Table for VRF "default"

(10.10.10.11/32, 227.0.0.1/32), uptime: 00:18:21, static ip pim
  Incoming interface: Ethernet1/1, RPF nbr: 10.10.10.11
  Outgoing interface list: (count: 2)
    Vlan200, uptime: 00:15:48, static
    Ethernet1/2, uptime: 00:18:21, static
```

MFDM - Multicast FIB Distribution Manager is responsible for distributing the multicast update messages (get from MRIB) to all the relevant line cards and the standby supervisor.

show forwarding distribution multicast route

```
(10.10.10.11/32, 227.0.0.1/32), RPF Interface: Ethernet1/1, flags:
Received Packets: 0 Bytes: 0
Number of Outgoing Interfaces: 2
Outgoing Interface List Index: 19
Vlan200
Ethernet1/2
```

show forwarding distribution multicast outgoing-interface-list L3 19

```
Outgoing Interface List Index: 19
Reference Count: 1
Platform Index: 0x2835
Number of Outgoing Interfaces: 2
Vlan200
Ethernet1/2
```

MFIB - (Multicast Forwarding Information Base) is used to route the packet on the linecard, it also contain hardware-specific information on how to replicate the packet across line cards.

Check the (VPN, S, G) programming in FIB-PI

module-1# show forwarding multicast route

```
IPv4 Multicast Routing table table-id:1
Total number of groups: 2
Legend:
C = Control Route
D = Drop Route
G = Local Group (directly connected receivers)
O = Drop on RPF failure
P = Punt to Supervisor
W = Wildcard
d = Decap route
N = VPC Non-Forwarder

(10.10.10.11/32, 227.0.0.1/32), RPF Interface: Ethernet1/1, flags:
Received Packets: 0 Bytes: 0
Number of Outgoing Interfaces: 2
Outgoing Interface List Index: 19
Vlan200 Outgoing Packets:0 Bytes:0
Ethernet1/2 Outgoing Packets:0 Bytes:0
```

module-1# show forwarding multicast outgoing-interface-list

```
Outgoing Interface List Index: 19
Reference Count: 1
```

```
Vlan200
Ethernet1/2
```

Check the (VPN, S, G) programming in FIB-PD to make sure that the routes are correctly programmed

module-1# show system internal forwarding multicast route detail

Hardware Multicast FIB Entries:

Flags Legend:

```
* - s_star_priority
S - sg_entry
D - Non-RPF Drop
B - Bi-dir route  W - Wildcard route
```

```
(10.10.10.11/32, 227.0.0.1/32), Flags: *S
Dev: 0, HWIndex: 0x1039be DRAM Index:: 0x1039be, VPN: 0x1
RPF Interface: Ethernet1/1, LIF: 0x1000
ML3 Adj/Rit Idx: 0x8d/0x14a, INGRESS_MET: 0x5, EGRESS_MET: 0x3
PD oiflist Idx: 0xf
MD Adj/Rit Idx: 0x9c/0x1a4, MDT Idx: 0x2, MTU Idx: 0x1, Dest Idx: 0x2835
Dev: 0 Adj/Rit Index: 0x99/0x182
    Type: OIF      elif: 0x1001      Ethernet1/2
    Dest Idx: 0xbfe      SMAC: 547f.eeeb.8bc1

(10.10.10.11/32, 227.0.0.1/32), Flags: *S
Dev: 1, HWIndex: 0x1039be DRAM Index:: 0x1039be, VPN: 0x1
RPF Interface: Ethernet1/1, LIF: 0x1000
ML3 Adj/Rit Idx: 0x8d/0x14a, INGRESS_MET: 0x4, EGRESS_MET: 0x2
PD oiflist Idx: 0xf
MD Adj/Rit Idx: 0x9c/0x1a4, MDT Idx: 0x2, MTU Idx: 0x1, Dest Idx: 0x2835
Dev: 1 Adj/Rit Index: 0x9b/0x1a2
    Type: OIF      elif: 0xc8      Vlan200
    Dest Idx: 0x0      SMAC: 547f.eeeb.8bc1
```

Make sure that the FIB DRAM, ml3 Adjacency, RIT and MET list are programmed correctly for the route each instance.

**module-1# debug forwarding spl fib-tcam inst 0 lookup ipv4-mcast ipv4_da 227.0.0.1
ipv4_sa 10.10.10.11 vpn 1**

Use the Dev and VPN idx from the output above

```
index: 0x1039be priority: 0
```

module-1# show hardware internal forwarding l3 inst 0 table fib_rslt_mcast index 0x1039be

Use the HWIndex from the Hardware Multicast FIB bove

```
entry_type: 1
adj_ptr: 0x14a The adjacency Pointer on M3 is actually the RIT index
rpf_lif1: 4096 Ethernet 1/1
sg_entry: 1
s_star_priority: 1
```

module-1# show hardware internal forwarding l3 inst 0 table rit index 0x14a detail

Use the RIT idx from the output above

```
<-snip->
suppress_ml2: 0
clr_mark_en: 0
met0_ptr: 5 This is the ingress MET pointer of Dev 0
met1_ptr: 3 This is the egress MET pointer of Dev 0
md_vld: 0
opt_shim: 0
<-snip->
```

module-1# show hardware internal forwarding l3 inst 0 table rw2adj_map_tbl_adjptr index 0x14a

Use the RIT idx from the output above

```
adjptr: 141 Hex 0x8d, is the ml3 adjancency
```

module-1# show hardware internal rewrite_engine inst 0 table brp-met-tbl

Use this command to dump the Multicast Expansion Table (MET)

```
+-----+
| MET table (logical layout) for F4 Bridge
| Inst 0; port(s) 1-24
|
|
Only non-zero entries are shown. All values in Hex.
```

	IDX	V	J	E	S	ADJ	ML3	BL2	PRS	NISH	LOPI	LOPI	PC	UUC	BS	PLH	DLB	VSL	PEER	BYP
									RBH		VAL	SEL	IDX	SUP	SUP			DROP	LINK	IFE
(0000)	1	0	1	0	00000	1	0	0	0	0000	0	000	0	0	0	0	0	0	0	1
(0003)	1	0	0	0	00182	1	0	0	0	1001	1									

```

fff  0  0  0  0  0  0  0  1  Egress RIT idx 0x182 - oif rit index, 0x1001 -
LIF value for eth 1/2
(0004) 1 0 1 0 0000 1 0 0 0 0000 0 000 0 0 0 0 0 0 1
(0005) 1 0 1 0 001a4 1 1 1 0 0000 1
000  0  0  1  0  0  0  0  1  Ingress RIT idx 0x1a4 - MD rit index

```

module-1# show system internal iftmc hardware lif brief | grep 1001

Use this command to map the LIF value to its logical interface

```
0x1a001000 0x4003 0x1001 0x1001
```

module-1# show system internal iftmc info interface brief | grep 0x1a001000

```
Eth1/2      0x1a001000 0x4003 0xbfe INTF LIF      UP      L3
```

Make sure that the MD Adjacency, RIT and MDT are programmed correctly for MD .

MD is used to forward packets to be sent to other line cards for egress replication.

module-1# show hardware internal forwarding l3 inst 0 table rw2adj_map_tbl_adjptr index 0x1a4

Use the MD adj RIT idx from Hardware Multicast FIB output.

```
adjptr: 156 hex 0x9c
```

module-1# show hardware internal forwarding l3 inst 0 table adj index 0x9c

Use the adj ptr from the above output.

```

<-snip->
no_intra_split_horizon 0x0
egress_lif 0x2
ri 0x0
top_sel 0x0
zone_enforce_en_or_use_vft 0x0
filter_en 0x0
<-snip->

```

module-1# show hardware internal forwarding l3 inst 0 table mdt index 0x2

Use the egress lif from the above output.

```
dest_idx: 10293  hex 0x2835
mtu_index: 1
```

Check the programming on the ingress LTL table MD-DI

module-1# show hardware internal rewrite_engine inst 0 table bpl-nec-tbl index 0x2835

Use the destination idx from the above output.

```
+-----+
| NEC table (logical layout) for F4 Bridge
| Inst 0; port(s) 1-24
|
```

Only non-zero entries are shown

IDX	RAW	EC VLD	MAP SEL	DROP	VQI VLD	VQI/ MI/ ECADDR	MI EGR	CS	LS	HBP
(0x2835)	0x000000000000800	0	0x00	0	0	0x00001	0x0000	0	0	0

Will give you inst for the fabric port.

module-1# show hardware internal fabric local inst 1 driver_info

Use the inst from the above output.

```
+-----+
| Instance Data Structures for SM15 Xbar ASIC
| Inst 0
```

```
inst number ..... 1
asic hw verion ..... 2
port bitmap ..... 0x03ffd360
port string ..... 5-6,8-9,12,14-25
slot ..... 0
power state ..... enabled
fabric present ..... yes
online ..... yes
serdes image ..... /lc/isan/bin/fencer_serdes.rom
first init done ..... yes
failure ..... no
hwaccess error ..... no
```

```
-----+-----+-----+-----+-----+
Port-Enabled    Connected-To          Mode          Chan-Status    Speed
-----+-----+-----+-----+-----+
```

05	Xbar-IF ioslice-02 link-03	fencer	Synced	116 Gbps
06	Xbar-IF ioslice-02 link-04	fencer	Synced	116 Gbps
08	Xbar-IF ioslice-01 link-03	fencer	Synced	116 Gbps
09	Xbar-IF ioslice-01 link-04	fencer	Synced	116 Gbps
22	Fab-2 link-01	chico	Synced	55 Gbps
23	Fab-1 link-01	chico	Synced	55 Gbps
24	Fab-2 link-02	chico	Synced	55 Gbps
25	Fab-1 link-02	chico	Synced	55 Gbps

module-1# show hardware internal fabric local inst 1 port 22 fpoe mcast

Use the FPOE from the above output

```
Fabric 0 Inst 1 Port 22 Multicast FPOE Database
index (hex)      value      ports
-----
10293 (2835)     0x00000120    5,8    MD packet from ingress will be sent to slot 5
and 8
49132 (bfec)     0x00000020     5
49151 (bfff)     0x00000120    5,8
49178 (c01a)     0x00000020     5
```

Make sure that the Adjacency, RIT and ELM are programmed correctly for the outgoing interface.

module-1# show hardware internal forwarding l3 inst 0 table rit index 0x182

Use the egress RIT from the MET table dump

```
d: 1
ccc: 4
smac: 1
l2_smac_ptr: 1    Check the smac table to see if the smac is programmed correctly
iptomac: 1
```

module-1# show hardware internal forwarding l2 inst 0 table smac_tbl index 0x1

Use the l2_smac_ptr from the above output

```
+-----+
| Source MAC table (logical layout) for F4 L2FWD driver
| Inst 0; port(s) 1-24
|
```

INDEX	VALUE
-------	-------

DEC	/	HEX	HEX
1	/	1	0000547f eeeb8bc1

module-1# show hardware internal forwarding l3 inst 0 table rw2adj_map_tbl_adjptr index 0x182

Use the egress RIT from the MET table dump

adjptr: 153 hex 0x99

module-1# show hardware internal forwarding l3 inst 0 table adj index 0x99

Use the egress adj ptr from the above output.

```

same_if_mask_sel 0x0
ingress_lif_segid_sel 0x0
format 0x0
fc_iod_drop 0x0
mcast_cpp_lif 0x0
ad_age 0x0
l3_enable 0x0
ad_trig 0x0
valid 0x1
rdt 0x1
peer_id_sel 0x0
no_intra_split_horizon 0x0
egress_lif 0x1001 LIF value for int eth 1/2
ri 0x0
top_sel 0x0
zone_enforce_en_or_use_vft 0x0
filter_en 0x1
frr_te 0x0
usd_da 0x0
gleen_adj 0x0
index_sel_or_bndl_en 0x1
tnl_encap 0x0
rw_hint 0x0
preserve_cos 0x0
ttl_control 0x2

```